

WE CLAIM:

1. A computer system, comprising:

5 a first cluster including a first plurality of processors and a first interconnection controller, the first plurality of processors and the first interconnection controller interconnected by first point-to-point intra-cluster links; and

10 a second cluster including a second plurality of processors and a second interconnection controller, the second plurality of processors and the second interconnection controller interconnected by second point-to-point intra-cluster links, the first interconnection controller coupled to the second interconnection controller by point-to-point inter-cluster links,

the first and second interconnection controllers configured to:

15 perform an initialization sequence that establishes a characteristic skew pattern between data lanes of the point-to-point inter-cluster links;

encode clock data in each symbol transmitted on the point-to-point inter-cluster links;

recover clock data from each symbol received on the point-to-point inter-cluster links; and

20 apply the characteristic skew pattern to correct for skew between data lanes of the point-to-point inter-cluster links.

2. The computer system of claim 1, wherein the encoding step comprises encoding 8-bit symbols as 10-bit symbols.

25 3. The computer system of claim 1, wherein the encoding step comprises encoding 4-bit symbols as 5-bit symbols.

4. The computer system of claim 1, wherein the initialization sequence comprises the use of one or more training sequences having known structures and lengths.

5 5. The computer system of claim 1, wherein the initialization sequence comprises establishing a phase lock loop for incoming frequency-encoded data.

6. An interconnection controller, comprising:

an intra-cluster interface configured for coupling with intra-cluster links to a plurality of local processors arranged in a point-to-point architecture in a local cluster;

an inter-cluster interface configured for coupling with an inter-cluster link to a non-local interconnection controller in a non-local cluster; and

a transceiver configured to:

15 perform an initialization sequence with the non-local interconnection controller that establishes a characteristic skew pattern between data lanes of the point-to-point inter-cluster links;

recover clock data from symbols received on the point-to-point inter-cluster links; and

20 apply the characteristic skew pattern to correct for skew between data lanes of the point-to-point inter-cluster links.

7. The interconnection controller of claim 6, further configured to encode clock data in symbols transmitted on the point-to-point inter-cluster links.

8. The interconnection controller of claim 6, further configured to forward symbols to the intra-cluster interface for transmission via the intra-cluster links to at least one of the plurality of local processors.

5 9. An integrated circuit comprising the interconnection controller of claim 6.

10. A set of semiconductor processing masks representative of at least a portion of the interconnection controller of claim 6.

10 11. At least one computer-readable medium having data structures stored therein representative of the interconnection controller of claim 6.

12. The integrated circuit of claim 9, wherein the integrated circuit comprises an application-specific integrated circuit.

15

13. The at least one computer-readable medium of claim 11, wherein the data structures comprise a simulatable representation of the interconnection controller.

20 14. The at least one computer-readable medium of claim 11, wherein the data structures comprise a code description of the interconnection controller.

15. The at least one computer-readable medium of claim 13, wherein the simulatable representation comprises a netlist.

25 16. The at least one computer-readable medium of claim 14, wherein the code description corresponds to a hardware description language.